

Election Verifiability: Cryptographic Definitions and an Analysis of Helios and JCJ

Ben Smyth

Huawei Technologies Co. Ltd.,
Boulogne-Billancourt, France
research@bensmyth.com

Steven Frink

Cornell University
Ithaca, NY, US
sfrink@cs.cornell.edu

Michael R. Clarkson

Cornell University
Ithaca, NY, US
clarkson@cs.cornell.edu

Abstract—Definitions of election verifiability in the computational model of cryptography are proposed. The definitions formalize notions of voters verifying their own votes, auditors verifying the tally of votes, and auditors verifying that only eligible voters vote. The Helios (Adida et al., 2009) and JCJ (Juels et al., 2010) election schemes are analyzed using these definitions. Helios 4.0 satisfies the definitions, but Helios 2.0 does not because of previously known attacks. JCJ does not satisfy the definitions because of a trust assumption it makes, but it does satisfy a weakened definition. Two previous definitions of verifiability (Juels et al., 2010; Cortier et al., 2014) are shown to permit election schemes vulnerable to attacks, whereas the new definitions prohibit those schemes.

I. INTRODUCTION

Electronic voting systems that have been deployed in real-world, large-scale public elections place extensive trust in software and hardware. Unfortunately, instead of being trustworthy, many systems are vulnerable to attacks that could bring election outcomes into disrepute [19], [45], [56], [84]. So relying solely on trust in voting systems is unwise; verification of election outcomes is essential.¹

Election verifiability enables voters and auditors to ascertain the correctness of election outcomes, regardless of whether the software and hardware of the voting system are trustworthy [2], [3], [25], [57], [75]. Kremer et al. [63] decompose election verifiability into three aspects:²

- *Individual verifiability*: voters can check that their own ballots are recorded.
- *Universal verifiability*: anyone can check that the tally of recorded ballots is computed properly.
- *Eligibility verifiability*: anyone can check that each tallied vote was cast by an authorized voter.

We propose new definitions of these three aspects of verifiability in the computational model of cryptography. We show that individual and universal verifiability are orthogonal, and that eligibility verifiability implies individual verifiability.

Because some electronic voting systems implement voter authentication themselves, whereas other systems outsource

voter authentication to third parties, we develop two variants of our definitions—one for systems with *internal authentication* and another for systems with *external authentication*. We employ our definitions to analyze the verifiability of two well-known election schemes, JCJ [59] and Helios [6]. JCJ is an election scheme that achieves *coercion resistance* and has been implemented as Civitas [29]; it implements its own internal authentication. Helios is a web-based voting system that has been deployed in the real-world and outsources authentication.

The Helios 2.0 election scheme is known to have vulnerabilities that enable attacks on verifiability, and several patches for those vulnerabilities have been proposed [17], [18], [33], [34]. By employing those proposed patches, we obtain a scheme called Helios 4.0 that satisfies our definition of election verifiability with external authentication. Helios 2.0, as expected, fails to satisfy our definition.

The JCJ election scheme does not satisfy our definition of eligibility verifiability, because an adversary who learns the taller's private key could cast unauthorized votes. We introduce a weakened definition of eligibility verifiability, incorporating JCJ's trust assumption that the private key is unknown to the adversary, and show that JCJ satisfies our weakened definition of election verifiability with internal authentication.

Our definitions of election verifiability improve upon two previous definitions [32], [59] by detecting a new class of *collusion attacks*, in which the tallying algorithm announces an incorrect tally, and the verification algorithm colludes with the tallying algorithm to accept the incorrect tally. Examples of collusion attacks include vote stuffing, and announcing tallies that are independent of the election. Our definitions also improve upon those previous definitions by detecting a new class of *biasing attacks*, in which the verification algorithm rejects some legitimate election outcomes. Examples of biasing attacks include rejecting outcomes in which a particular candidate does not win, and rejecting all election outcomes, even correct outcomes.

This paper thus contributes to the security of electronic voting systems by

- proposing computational definitions of election verifiabil-

¹Doveryai, no proveryai (trust, but verify) says the Russian proverb.

²This decomposition has been criticized [69]; we refute that criticism in Section VII.

- showing that individual, universal, and eligibility verifiability are mostly orthogonal properties of voting systems,
- proving that well-known election schemes do (or do not) satisfy election verifiability, and
- identifying collusion and biasing attacks as new classes of attacks on voting systems and demonstrating that they are not detected by two earlier definitions.

Ours are the first proofs that Helios 4.0 and JCJ satisfy a computational definition of verifiability.

Structure: Section II defines election verifiability with external authentication. Section III analyzes Helios. Section IV defines election verifiability with internal authentication. Section V analyzes JCJ. Section VI introduces collusion and biasing attacks. Section VII reviews related work, and Section VIII concludes.

II. EXTERNAL AUTHENTICATION

Some election schemes do not implement authentication themselves, but instead rely on an external authentication mechanism. Helios, for example, supports authentication with Facebook, Google and Yahoo credentials.³ In essence, the election scheme outsources ballot authentication. We begin by defining election verifiability for that model.

A. Election scheme

An *election scheme with external authentication*, which henceforth in this section we abbreviate as “election scheme,” is a tuple $(\text{Setup}, \text{Vote}, \text{Tally}, \text{Verify})$ of probabilistic polynomial-time (PPT) algorithms:

- **Setup**, denoted⁴ $(PK_{\mathcal{T}}, SK_{\mathcal{T}}, m_B, m_C) \leftarrow \text{Setup}(k)$, is executed by the *tallier*, who is responsible for tallying ballots.⁵ Setup takes a security parameter k as input and outputs a key pair $(PK_{\mathcal{T}}, SK_{\mathcal{T}})$, a maximum number of ballots m_B , and a maximum number of candidates m_C .
- **Vote**, denoted $b \leftarrow \text{Vote}(PK_{\mathcal{T}}, n_C, \beta, k)$, is executed by voters. A voter makes a *choice* of candidate from a sequence c_1, \dots, c_{n_C} of candidates. A *well-formed* choice is an integer β , such that $1 \leq \beta \leq n_C$. Vote takes as input the public key $PK_{\mathcal{T}}$ of the tallier, the number n_C of candidates, the voter’s choice β of candidate, and security parameter k . It outputs a ballot b , or error symbol \perp . An error might occur if the candidate choice is not well-formed or for other reasons particular to the election scheme.
- **Tally**, denoted $(\mathbf{X}, P) \leftarrow \text{Tally}(PK_{\mathcal{T}}, SK_{\mathcal{T}}, BB, n_C, k)$, is executed by the tallier. It involves a public *bulletin*

³https://github.com/benadida/helios-server/tree/master/helios_auth/auth_systems, accessed 4 Aug 2015.

⁴Let $\text{Alg}(\text{in}; r)$ denote the output of probabilistic algorithm Alg on input in and random coins r . Let $\text{Alg}(\text{in})$ denote $\text{Alg}(\text{in}; r)$, where r is chosen uniformly at random. And let \leftarrow denote assignment.

⁵Some election schemes (e.g., Helios and JCJ) permit the tallier’s role to be distributed amongst several talliers. For simplicity, we consider only a single tallier in this paper.

board BB , which we model as a set.⁶ Tally takes as input the public key $PK_{\mathcal{T}}$ and private key $SK_{\mathcal{T}}$ of the tallier, the bulletin board BB , the number of candidates n_C , and security parameter k . It outputs a tally \mathbf{X} and a non-interactive proof P that the tally is correct. A *tally* is a vector \mathbf{X} of length n_C such that $\mathbf{X}[j]$ indicates the number of votes for candidate c_j .⁷

- **Verify**, denoted $v \leftarrow \text{Verify}(PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P, k)$, can be executed by anyone to audit the election. Verify takes as input the public key $PK_{\mathcal{T}}$ of the tallier, the bulletin board BB , the number of candidates n_C , a tally \mathbf{X} , a proof P of correct tallying, and security parameter k . It outputs a bit v , which is 1 if the tally successfully verifies and 0 otherwise. We assume that Verify is deterministic.

Election schemes must satisfy Correctness, which asserts that tallies produced by Tally corresponds to the choices input to Vote:

Definition 1 (Correctness). *There exists a negligible function μ , such that for all security parameters k , integers n_B and n_C , and choices $\beta_1, \dots, \beta_{n_B} \in \{1, \dots, n_C\}$, it holds that if \mathbf{Y} is a vector of length n_C whose components are all 0, then*

$$\Pr[(PK_{\mathcal{T}}, SK_{\mathcal{T}}, m_B, m_C) \leftarrow \text{Setup}(k); \\ \text{for } 1 \leq i \leq n_B \text{ do} \\ \quad b_i \leftarrow \text{Vote}(PK_{\mathcal{T}}, n_C, \beta_i, k); \\ \quad \mathbf{Y}[\beta_i] \leftarrow \mathbf{Y}[\beta_i] + 1; \\ \quad BB \leftarrow \{b_1, \dots, b_{n_B}\}; \\ (\mathbf{X}, P) \leftarrow \text{Tally}(PK_{\mathcal{T}}, SK_{\mathcal{T}}, BB, n_C, k) : \\ n_B \leq m_B \wedge n_C \leq m_C \Rightarrow \mathbf{X} = \mathbf{Y}] > 1 - \mu(k).$$

Note that Correctness does not involve an adversary. Correctness therefore stipulates that, under ideal conditions, an election scheme does indeed produce the correct tally. Correctness is not actually necessary to achieve verifiability: our definition of universal verifiability will ensure that, in the presence of an adversary, Verify detects any errors in the tally. But it is reasonable to rule out election schemes that simply do not work properly under ideal conditions.

Election schemes must also satisfy Completeness, which stipulates that tallies produced by Tally will actually be accepted by Verify:

Definition 2 (Completeness). *There exists a negligible function μ , such that for all security parameters k , bulletin boards BB , and integers n_C , it holds that*

$$\Pr[(PK_{\mathcal{T}}, SK_{\mathcal{T}}, m_B, m_C) \leftarrow \text{Setup}(k); \\ (\mathbf{X}, P) \leftarrow \text{Tally}(PK_{\mathcal{T}}, SK_{\mathcal{T}}, BB, n_C, k) : \\ |BB| \leq m_B \wedge n_C \leq m_C \Rightarrow \\ \text{Verify}(PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P, k) = 1] > 1 - \mu(k).$$

⁶Bulletin boards have also been modeled as public broadcast channels [35], [76], [78]. We abstract from the details of channels by employing sets to represent the data sent on them. We favor sets over multisets, because Cortier and Smyth [33], [34] demonstrate attacks against privacy when the bulletin board is modeled as a multiset.

⁷Let $\mathbf{X}[i]$ denote component i of vector \mathbf{X} .

Without Completeness, election schemes might be vulnerable to biasing attacks, as we show in Section VI-B.

Finally, election schemes must satisfy Injectivity, which asserts that a ballot cannot be interpreted as a vote for more than one candidate:

Definition 3 (Injectivity). *For all security parameters k , public keys PK_T , integers n_C , and choices β and β' , such that $\beta \neq \beta'$, we have*

$$\Pr[b \leftarrow \text{Vote}(PK_T, n_C, \beta, k); \\ b' \leftarrow \text{Vote}(PK_T, n_C, \beta', k) : \\ b \neq \perp \wedge b' \neq \perp \Rightarrow b \neq b'] = 1.$$

Injectivity ensures that distinct choices are not mapped by Vote to the same ballot. Without Injectivity, an election scheme might produce ballots whose meaning is ambiguous. For example, if $\text{Vote}(PK_T, n_C, \beta, k; r)$ were defined to be $\beta + r$, then a ballot b could be tallied as any well-formed choice β' such that $\beta' = b - r'$ for some r' . But that definition of Vote is prohibited by Injectivity. Thus, Injectivity helps to ensure that the choices used to construct ballots can be uniquely tallied.

Limitations: Our model of election schemes is sufficient to analyze Helios and, after we extend the model to handle internal authentication in Section IV-A, JCJ. These are notable schemes, and formally analyzing their verifiability is a novel contribution. But there are other notable schemes that fall outside our model:

- Pret à Voter [25], MarkPledge [73], Scantegrity II [24], and Remotegrity [85] all rely on features implemented with paper, such as scratch-off surfaces and detachable columns.
- Everlasting privacy [72], which requires Vote to output a public ballot and a secret proof, involving temporal information, to the voter.
- Scytł's Pnyx.core ODBP 1.0 [28], which requires the bulletin board to be divided into two parts: a public part visible to all participants, and a secret part visible only to election administrators.

We leave extension of our model to other election schemes as future work.

B. Election verifiability

Election verifiability comprises three aspects: individual, universal, and eligibility verifiability. We express each as an *experiment*, which is an algorithm that outputs 0 or 1. The adversary *wins* an experiment by causing it to output 1.

1) *Individual verifiability:* In our model of election schemes, all recorded ballots are posted on the bulletin board. So for a voter to verify that their ballot has been recorded, it suffices to enable them to uniquely identify their ballot on the bulletin board.⁸

Individual verifiability experiment $\text{Exp-IV-Ext}(\Pi, \mathcal{A}, k)$, where Π denotes an election scheme, \mathcal{A} denotes the adversary,

⁸Section VIII addresses the complementary issue of whether a recorded ballot corresponds to the candidate choice a voter intended to make.

and k denotes a security parameter, therefore challenges \mathcal{A} to generate a scenario in which the voter cannot uniquely identify their ballot. In essence, Exp-IV-Ext challenges \mathcal{A} to generate a collision from Vote.⁹ If \mathcal{A} cannot win, then voters can uniquely identify their ballots on the bulletin board:

```
Exp-IV-Ext( $\Pi, \mathcal{A}, k$ ) =
1  $(PK_T, n_C, \beta, \beta') \leftarrow \mathcal{A}(k);$ 
2  $b \leftarrow \text{Vote}(PK_T, n_C, \beta, k);$ 
3  $b' \leftarrow \text{Vote}(PK_T, n_C, \beta', k);$ 
4 if  $b = b' \wedge b \neq \perp \wedge b' \neq \perp$  then
5   return 1
6 else
7   return 0
```

Line 1 asks \mathcal{A} to compute two candidate choices β and β' , such that ballots b and b' for those choices, as computed by Vote in lines 2 and 3, are equal. Individual verifiability thus resembles Injectivity, but individual verifiability allows choices to be equal and allows \mathcal{A} to choose election parameters.

One way to achieve individual verifiability is to base the election scheme on a probabilistic encryption scheme, such as El Gamal [41]. Intuitively, if Vote encrypts the choice using random coins, then it is overwhelmingly unlikely that two votes will result in the same ballot. Our proofs that Helios and JCJ satisfy individual verifiability are based on this idea.

Clash attacks: In a *clash attack* [71], the adversary convinces some voters that a single ballot belongs to them all. Some clash attacks are possible because of vulnerabilities in the design of Vote. For example, if Vote simply outputs candidate choice β , then a voter has no way to distinguish their vote for β from another voter's vote for β . Exp-IV-Ext detects clash attacks resulting from vulnerabilities in Vote.

Some clash attacks, however, are possible because the adversary subverts the implementation of Vote. For example, the adversary might replace some hardware or software, or compromise the random number generator. If any one of these aspects is compromised, then Vote has effectively been changed to a different algorithm Vote' . The conclusions drawn by a security analyst who uses our definition of individual verifiability to analyze Vote would not necessarily be applicable to Vote' .

In short, a voter can verify that their ballot has been recorded if and only if they run the correct Vote algorithm. We make no guarantees to voters that do not run the correct Vote algorithm. One way to make stronger guarantees is to use cut-and-choose protocols to audit ballots [11], [12]. This would require modeling voting as an interactive protocol with the adversary, rather than as an algorithm. We leave this extension as future work.

2) *Universal verifiability:* For an election to be universally verifiable, anyone must be able to check that a tally is correct with respect to recorded ballots—that is, the tally represents

⁹Exp-IV-Ext can be equivalently formulated as an experiment that challenges \mathcal{A} to predict the output of Vote. See the companion technical report [1] for details.

the choices used to construct the recorded ballots. Because anyone can execute Verify, it suffices that Verify accepts only when that property holds.

Universal verifiability experiment $\text{Exp-UV-Ext}(\Pi, \mathcal{A}, k)$ therefore challenges adversary \mathcal{A} to concoct a scenario in which Verify incorrectly accepts:

```
 $\text{Exp-UV-Ext}(\Pi, \mathcal{A}, k) =$ 
1  $(PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P) \leftarrow \mathcal{A}(k);$ 
2  $\mathbf{Y} \leftarrow \text{correct-tally}(PK_{\mathcal{T}}, BB, n_C, k);$ 
3  $\text{if } \mathbf{X} \neq \mathbf{Y} \wedge \text{Verify}(PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P, k) = 1 \text{ then}$ 
4    $\quad \text{return } 1$ 
5  $\text{else}$ 
6    $\quad \text{return } 0$ 
```

In line 1, \mathcal{A} is challenged to create a bulletin board BB and purported tally \mathbf{X} of that bulletin board. Line 2 constructs the correct tally \mathbf{Y} of BB , and line 3 checks whether Verify accepts an incorrect tally. If \mathcal{A} cannot win Exp-UV-Ext , then Verify will not accept incorrect tallies. In particular, no ballots can be omitted from the tally, and at most one candidate choice can be included in the tally for each ballot.

Let function correct-tally be defined such that for all $PK_{\mathcal{T}}$, BB , n_C , k , ℓ , and $\beta \in \{1, \dots, n_C\}$,

$$\begin{aligned} \text{correct-tally}(PK_{\mathcal{T}}, BB, n_C, k)[\beta] &= \ell \\ \iff \exists^{\leq \ell} b \in (BB \setminus \{\perp\}) : \\ \exists r : b &= \text{Vote}(PK_{\mathcal{T}}, n_C, \beta, k; r). \end{aligned}$$

The vector produced by correct-tally must be of length n_C . Component β of vector $\text{correct-tally}(PK_{\mathcal{T}}, BB, n_C, k)$ equals ℓ iff there exist¹⁰ ℓ ballots on the bulletin board that are votes for candidate β . It follows that the output of correct-tally represents the choices used to construct the recorded ballots. Note that, without Injectivity, the existential quantification in correct-tally could permit a ballot to be tallied for more than one candidate. Of course, correct-tally cannot be computed by a PPT algorithm for typical cryptographic election schemes. But that does not matter, because correct-tally is never actually computed as part of an election scheme—its use is solely in the definition of Exp-UV-Ext .¹¹

Security analysts must convince themselves that correct-tally is indeed correct. Because of the function's simplicity, this should be relatively straightforward. By comparison, Tally algorithms for real voting schemes tend to be complicated. For example, compare the complexity of correct-tally to Helios's Tally algorithm, which appears in the companion technical report [1].

By design, Exp-UV-Ext assumes that the ballots on bulletin board BB are exactly the ballots that should be tallied. The external authentication mechanism is assumed to prohibit

¹⁰The definition of correct-tally employs a *counting quantifier* [79] denoted $\exists^=$. Predicate $(\exists^= \ell x : P(x))$ holds exactly when there are ℓ distinct values for x such that $P(x)$ is satisfied. Variable x is bound by the quantifier, whereas ℓ is free.

¹¹Kiayias et al. [61] use a similar super-polynomial *vote extractor* to recover choices from ballots in an experiment defining verifiability.

unauthorized ballots from being posted on BB . Helios makes such an assumption about its external authentication mechanism.

3) *Eligibility verifiability*: For an election to satisfy eligibility verifiability, anyone must be able to check that every tallied vote was cast by an authorized voter—that is, it must be possible to authenticate ballots. In election schemes with external authentication, a trusted third party authenticates ballots. That third party might convince itself that all tallied ballots have been authenticated, but it cannot convince all other parties. Eligibility verifiability, therefore, is not achievable in election schemes with external authentication.

4) *Election verifiability*: With Exp-IV-Ext and Exp-UV-Ext , we define election verifiability with external authentication. Let a PPT adversary's *success* $\text{Succ}(\text{Exp}(\cdot))$ in an experiment $\text{Exp}(\cdot)$ be the probability that the adversary wins—that is, $\text{Succ}(\text{Exp}(\cdot)) = \Pr[\text{Exp}(\cdot) = 1]$.

Definition 4 (Ver-Ext). *An election scheme Π satisfies election verifiability with external authentication (Ver-Ext) if for all PPT adversaries \mathcal{A} , there exists a negligible function μ , such that for all security parameters k , it holds that $\text{Succ}(\text{Exp-IV-Ext}(\Pi, \mathcal{A}, k)) + \text{Succ}(\text{Exp-UV-Ext}(\Pi, \mathcal{A}, k)) \leq \mu(k)$.*

An election scheme satisfies individual verifiability if $\text{Succ}(\text{Exp-IV-Ext}(\Pi, \mathcal{A}, k)) \leq \mu(k)$, and similarly for universal verifiability.

C. Example—Toy scheme from nonces

A toy election scheme satisfying Ver-Ext can be based on nonces. Each voter publishes a nonce paired with her choice of candidate to the bulletin board. This scheme illustrates the essence of election verifiability, even though it does not offer any privacy.

Definition 5. *Election scheme Nonce is defined as follows:*

- $\text{Setup}(k)$ outputs $(\perp, \perp, p_1(k), p_2(k))$, where p_1 and p_2 may be any polynomial functions.
- $\text{Vote}(PK_{\mathcal{T}}, n_C, \beta, k)$ selects a nonce r uniformly at random from \mathbb{Z}_{2^k} and outputs (r, β) .
- $\text{Tally}(PK_{\mathcal{T}}, SK_{\mathcal{T}}, BB, n_C, k)$ computes a vector \mathbf{X} of length n_C , such that \mathbf{X} is a tally of the votes on BB for which the nonce is in \mathbb{Z}_{2^k} , and outputs (\mathbf{X}, \perp) .
- $\text{Verify}(PK_{\mathcal{T}}, BB, n_C, \mathbf{X}, P, k)$ outputs 1 if $(\mathbf{X}, P) = \text{Tally}(\perp, \perp, BB, n_C, k)$ and 0 otherwise.

Proposition 1. *Nonce satisfies Ver-Ext.*

Proof sketch. Nonce satisfies individual verifiability, because voters can use their nonce to check that their own ballot appears on the bulletin board. With overwhelming probability, Vote will select unique nonces for each voter, hence generate distinct ballots. Nonce also satisfies universal verifiability, because plaintext candidate choices are posted on the bulletin board. \square

D. Orthogonality

Exp-IV-Ext and Exp-UV-Ext capture orthogonal security properties. A scheme that satisfies individual verifiability but violates universal verifiability can be constructed from Nonce by modifying Verify to always output 1. Voters can still check that their own ballot appears. But an adversary can easily win Exp-UV-Ext, because Verify will accept any tally. A scheme that satisfies universal verifiability but violates individual verifiability can be constructed from Nonce by removing the nonces, leaving just the voter’s choice in the ballots. Call that scheme Choice. Anyone can still verify the tally of the election, but an adversary can easily win Exp-IV-Ext, because two votes for the same candidate will collide.

III. CASE STUDY: HELIOS

Helios is an open-source, web-based electronic voting system.¹² Helios has been deployed in the real-world: the International Association of Cryptologic Research (IACR) has used Helios annually since 2010 to elect board members [14], [47], [53], the Catholic University of Louvain used Helios to elect the university president [6], and Princeton University has used Helios to elect several student governments [4], [74].

Attacks have been discovered against the original Helios scheme, and defenses against those attacks have been proposed [17], [18], [33], [34]. For clarity, we write *Helios 2.0* to refer to the Helios scheme as originally proposed [6] and *Helios 4.0* to refer to a version of Helios that incorporates the defenses.¹³ When referring in general to both of these schemes, we simply write *Helios*.

To achieve verifiability while maintaining *ballot secrecy* [16], [18], Helios homomorphically encrypts candidate choices. During tallying, all encrypted choices are homomorphically combined¹⁴ into a single ciphertext, which is then decrypted to reveal the tally. Informally, Helios works as follows:

- **Setup.** The tallier generates a key pair for a homomorphic encryption scheme and publishes the public key.
- **Voting.** A voter encrypts her candidate choice with the tallier’s public key, and she proves in zero knowledge that the ciphertext contains a well-formed choice. The voter posts her ballot (i.e., ciphertext and proof) on the bulletin board. During posting, the bulletin board is assumed to correctly authenticate voters.
- **Tallying.** The tallier discards any ballots from the bulletin board for which proofs do not hold. The tallier homomorphically combines the ciphertexts in the remaining ballots, decrypts the homomorphic combination, and proves

¹²<https://vote.heliosvoting.org/>

¹³Our analysis of Helios 4.0 is based on the specification [5] for the next release. This specification incorporates proposals by Cortier and Smyth [34] for non-malleable ballots and by Bernhard et al. [18] to replace the weak Fiat–Shamir transformation with the strong Fiat–Shamir transformation.

¹⁴The homomorphic combination of ciphertexts is straightforward for two-candidate elections [10], [15], [30], [50], [77], since choices (e.g., “yes” or “no”) can be encoded as 1 or 0. Multi-candidate elections are also possible [15], [37], [49].

in zero knowledge that decryption was performed correctly. Finally, the tallier publishes the winning candidate and proof of correct decryption.

- **Verification.** A verifier recomputes the homomorphic combination and checks all the zero-knowledge proofs.

We give a formal description of Helios 4.0 in the companion technical report [1].¹⁵ Using that formalization, we can prove that Helios 4.0 is verifiable:

Theorem 2. *Helios 4.0 satisfies Ver-Ext.*

Proof sketch. Helios 4.0 satisfies individual verifiability, because the probabilistic encryption scheme ensures that ballots are unique, with overwhelming probability. And Helios 4.0 satisfies universal verifiability, because the zero-knowledge proofs can be publicly verified. \square

A formal proof of Theorem 2 appears in the companion technical report [1]. The proof assumes the random oracle model [9].

We would not expect Ver-Ext to hold for Helios 2.0, because of known attacks [18]. Accordingly, we prove that Helios 2.0 does not satisfy Ver-Ext in the companion technical report [1].

IV. INTERNAL AUTHENTICATION

Some election schemes implement their own authentication mechanisms. JCJ [57]–[59] and Civitas [29], for example, authenticate ballots based on *credentials* issued to voters by a registration authority. Schemes with this kind of internal authentication enable verification of whether tallied ballots were cast by authorized voters.

A. Election scheme

A *registrar* is responsible for issuing authentication *credentials* to voters.¹⁶ Each voter is associated with a credential pair (pk, sk) . The voter uses private credential sk to construct a ballot. Public credential pk is used during tallying and verification. Let L denote the *electoral roll*, which is the set of all public credentials.

An *election scheme with internal authentication*, which henceforth in this section we abbreviate as “election scheme,” is a tuple $(\text{Setup}, \text{Register}, \text{Vote}, \text{Tally}, \text{Verify})$ of PPT algorithms. The algorithms are now denoted as follows:

- $(PK_{\mathcal{T}}, SK_{\mathcal{T}}, m_B, m_C) \leftarrow \text{Setup}(k)$
- $(pk, sk) \leftarrow \text{Register}(PK_{\mathcal{T}}, k)$
- $b \leftarrow \text{Vote}(sk, PK_{\mathcal{T}}, n_C, \beta, k)$
- $(\mathbf{X}, P) \leftarrow \text{Tally}(PK_{\mathcal{T}}, SK_{\mathcal{T}}, BB, L, n_C, k)$
- $v \leftarrow \text{Verify}(PK_{\mathcal{T}}, BB, L, n_C, \mathbf{X}, P, k)$

Setup is unchanged from election schemes with external authentication (cf. §II-A). The only change to *Vote* is that it now accepts private credential sk as input. Similarly, the only change to *Tally* and *Verify* is that they now accept electoral

¹⁵Our formalization is the first cryptographic description of Helios 4.0, hence an additional contribution of this work.

¹⁶Some election schemes (e.g., JCJ) permit the registrar’s role to be distributed among several registrars. For simplicity, we consider only a single registrar in this paper.

roll L as input. Register is executed by the registrar. It takes as input the public key $PK_{\mathcal{T}}$ of the tallier and security parameter k , and it outputs a *credential pair* (pk, sk) . After all voters have been registered, the registrar certifies the electoral roll, perhaps by digitally signing and publishing it.¹⁷

Election schemes must continue to satisfy Correctness, Completeness, and Injectivity, which we update to include private credentials and the electoral roll:

Definition 6 (Correctness). *There exists a negligible function μ , such that for all security parameters k , integers n_B and n_C , and choices $\beta_1, \dots, \beta_{n_B} \in \{1, \dots, n_C\}$, it holds that if \mathbf{Y} is a vector of length n_C whose components are all 0, then*

```

 $\Pr[(PK_{\mathcal{T}}, SK_{\mathcal{T}}, m_B, m_C) \leftarrow \text{Setup}(k);$ 
 $\quad \text{for } 1 \leq i \leq n_B \text{ do}$ 
 $\quad \quad (pk_i, sk_i) \leftarrow \text{Register}(PK_{\mathcal{T}}, k);$ 
 $\quad \quad b_i \leftarrow \text{Vote}(sk_i, PK_{\mathcal{T}}, n_C, \beta_i, k);$ 
 $\quad \quad \mathbf{Y}[\beta_i] \leftarrow \mathbf{Y}[\beta_i] + 1;$ 
 $\quad L \leftarrow \{pk_1, \dots, pk_{n_B}\};$ 
 $\quad BB \leftarrow \{b_1, \dots, b_{n_B}\};$ 
 $\quad (\mathbf{X}, P) \leftarrow \text{Tally}(PK_{\mathcal{T}}, SK_{\mathcal{T}}, BB, L, n_C, k) :$ 
 $\quad n_B \leq m_B \wedge n_C \leq m_C \Rightarrow \mathbf{X} = \mathbf{Y}] > 1 - \mu(k).$ 

```

Definition 7 (Completeness). *There exists a negligible function μ , such that for all security parameters k , bulletin boards BB , and integers n_C and n_V , it holds that*

```

 $\Pr[(PK_{\mathcal{T}}, SK_{\mathcal{T}}, m_B, m_C) \leftarrow \text{Setup}(k);$ 
 $\quad \text{for } 1 \leq i \leq n_V \text{ do } (pk_i, sk_i) \leftarrow \text{Register}(PK_{\mathcal{T}}, k);$ 
 $\quad L \leftarrow \{pk_1, \dots, pk_{n_V}\};$ 
 $\quad (\mathbf{X}, P) \leftarrow \text{Tally}(PK_{\mathcal{T}}, SK_{\mathcal{T}}, BB, L, n_C, k) :$ 
 $\quad |BB| \leq m_B \wedge n_C \leq m_C \Rightarrow$ 
 $\quad \text{Verify}(PK_{\mathcal{T}}, BB, L, n_C, \mathbf{X}, P, k) = 1] > 1 - \mu(k).$ 

```

Definition 8 (Injectivity). *For all security parameters k , public keys $PK_{\mathcal{T}}$, integers n_C , and choices β and β' , such that $\beta \neq \beta'$, we have*

```

 $\Pr[(pk, sk) \leftarrow \text{Register}(PK_{\mathcal{T}}, k);$ 
 $\quad (pk', sk') \leftarrow \text{Register}(PK_{\mathcal{T}}, k);$ 
 $\quad b \leftarrow \text{Vote}(sk, PK_{\mathcal{T}}, n_C, \beta, k);$ 
 $\quad b' \leftarrow \text{Vote}(sk', PK_{\mathcal{T}}, n_C, \beta', k) :$ 
 $\quad b \neq \perp \wedge b' \neq \perp \Rightarrow b \neq b'] = 1.$ 

```

B. Election verifiability

Recall (from §II-B) that election verifiability is expressed with experiments, and that an adversary wins by causing an experiment to output 1. We henceforth assume that the adversary is *stateful*—that is, information persists across invocations of the adversary in a single experiment. Our experiments in

¹⁷It might seem surprising that Register does not require the registrar to provide any private keys as input. But in constructions of election schemes with internal authentication, e.g., [29], [59], the registrar does not sign credential pairs with its own private key. Rather, the registrar signs the electoral roll.

Section II did not need this assumption, because they never invoked the adversary more than once.

In our experiments, below, we model an adversary who cannot corrupt the registration process that issues credentials to voters.¹⁸ Hence our definitions will not detect attacks against verifiability that result solely from weaknesses in the registration process. Secure construction of electoral rolls is not a topic that electronic voting usually addresses—though it seems an important part of any real-world deployment.

1) *Individual verifiability*: The individual verifiability experiment again challenges adversary \mathcal{A} to generate a scenario in which the voter could not uniquely identify their ballot.¹⁹

```

Exp-IV-Int( $\Pi, \mathcal{A}, k$ ) =
1  $(PK_{\mathcal{T}}, n_V) \leftarrow \mathcal{A}(k);$ 
2  $\text{for } 1 \leq i \leq n_V \text{ do } (pk_i, sk_i) \leftarrow \text{Register}(PK_{\mathcal{T}}, k)$ 
3  $L \leftarrow \{pk_1, \dots, pk_{n_V}\};$ 
4  $Crpt \leftarrow \emptyset;$ 
5  $(n_C, \beta, \beta', i, j) \leftarrow \mathcal{A}^C(L);$ 
6  $b \leftarrow \text{Vote}(sk_i, PK_{\mathcal{T}}, n_C, \beta, k);$ 
7  $b' \leftarrow \text{Vote}(sk_j, PK_{\mathcal{T}}, n_C, \beta', k);$ 
8  $\text{if}$ 
 $\quad b = b' \wedge b \neq \perp \wedge b' \neq \perp \wedge i \neq j \wedge sk_i \notin Crpt \wedge sk_j \notin Crpt$ 
 $\quad \text{then}$ 
9  $\quad \quad \text{return } 1$ 
10  $\text{else}$ 
11  $\quad \quad \text{return } 0$ 

```

The main differences from the corresponding experiment for external authentication (§II-B1) are that voters are registered in line 2, and that \mathcal{A} is given access to an oracle C in line 5. The oracle is used to model \mathcal{A} corrupting voters and learning their private credentials: on invocation $C(\ell)$, where $1 \leq \ell \leq n_V$, the oracle records that voter ℓ is corrupted by updating $Crpt$ to be $Crpt \cup \{sk_\ell\}$ and outputs sk_ℓ . In line 5, the voter indices output by \mathcal{A} must be legal with respect to n_V , but we elide that detail from the experiment for simplicity. Line 8 ensures that \mathcal{A} cannot trivially win by corrupting voters.

2) *Universal verifiability*: The universal verifiability experiment again challenges \mathcal{A} to concoct a scenario in which Verify incorrectly accepts:

¹⁸Küsters and Truderung [67] explore some consequences of permitting adversarial influence during registration.

¹⁹Unlike Exp-IV-Ext, a variant of Exp-IV-Int that challenges \mathcal{A} to predict the output of Vote is strictly stronger. See the companion technical report [1] for details.

```

Exp-UV-Int( $\Pi, \mathcal{A}, k$ ) =
1  $(PK_{\mathcal{T}}, n_V) \leftarrow \mathcal{A}(k)$ ;
2 for  $1 \leq i \leq n_V$  do  $(pk_i, sk_i) \leftarrow \text{Register}(PK_{\mathcal{T}}, k)$ 
3  $L \leftarrow \{pk_1, \dots, pk_{n_V}\}$ ;
4  $M \leftarrow \{(pk_1, sk_1), \dots, (pk_{n_V}, sk_{n_V})\}$ ;
5  $(BB, n_C, \mathbf{X}, P) \leftarrow \mathcal{A}(M)$ ;
6  $\mathbf{Y} \leftarrow \text{correct-tally}(PK_{\mathcal{T}}, BB, M, n_C, k)$ ;
7 if  $\mathbf{X} \neq \mathbf{Y} \wedge \text{Verify}(PK_{\mathcal{T}}, BB, L, n_C, \mathbf{X}, P, k) = 1$  then
8   return 1
9 else
10  return 0

```

The main differences from the corresponding experiment for external authentication (§II-B2) are that voters are registered in line 2, and their credential pairs are used in the rest of the experiment.

Function *correct-tally* is now modified to tally only authorized ballots. A ballot is *authorized* if it is constructed with a private credential from M , and that private credential was not used to construct any other ballot on BB . By comparison, the original *correct-tally* function (§II-B2) tallies all the ballots on BB .

Formally, let function *correct-tally* now be defined such that for all $PK_{\mathcal{T}}$, BB , M , n_C , k , ℓ , and $\beta \in \{1, \dots, n_C\}$,

$$\begin{aligned} \text{correct-tally}(PK_{\mathcal{T}}, BB, M, n_C, k)[\beta] &= \ell \\ \iff \exists^{=\ell} b \in \text{authorized}(PK_{\mathcal{T}}, (BB \setminus \{\perp\}), M, n_C, k) : \\ &\exists sk, r : b = \text{Vote}(sk, PK_{\mathcal{T}}, n_C, \beta, k; r). \end{aligned}$$

Let *authorized* be defined as follows:

$$\begin{aligned} \text{authorized}(PK_{\mathcal{T}}, BB, M, n_C, k) = \\ \{b : b \in BB \\ \wedge \exists pk, sk, \beta, r : b = \text{Vote}(sk, PK_{\mathcal{T}}, n_C, \beta, k; r) \\ \wedge (pk, sk) \in M \wedge \neg \exists b', \beta', r' : b' \in (BB \setminus \{b\}) \\ \wedge b' = \text{Vote}(sk, PK_{\mathcal{T}}, n_C, \beta', k; r')\}. \end{aligned}$$

Function *authorized* discards all revotes—that is, if there is more than one ballot submitted with a private credential sk , then all ballots submitted under that credential are discarded. Therefore, election schemes that permit revoting cannot be analyzed with this definition of *authorized*. But alternative definitions of *authorized* are possible—for example, if ballots were timestamped, *authorized* could discard all but the most recent ballot submitted under a particular credential.

3) *Eligibility verifiability*: Recall (from §II-B3) that for an election scheme to satisfy eligibility verifiability, anyone must be able to check that every tallied vote was cast by an authorized voter—that is, it must be possible to authenticate ballots. Because voters are issued credential pairs that can be used to authenticate ballots, it suffices to ensure that knowledge of a private credential is necessary to construct an authentic ballot.

Eligibility verifiability experiment Exp-EV-Int therefore challenges \mathcal{A} to produce a ballot under a private credential that \mathcal{A} does not know:

```

Exp-EV-Int( $\Pi, \mathcal{A}, k$ ) =
1  $(PK_{\mathcal{T}}, n_V) \leftarrow \mathcal{A}(k)$ ;
2 for  $1 \leq i \leq n_V$  do  $(pk_i, sk_i) \leftarrow \text{Register}(PK_{\mathcal{T}}, k)$ ;
3  $L \leftarrow \{pk_1, \dots, pk_{n_V}\}$ ;
4  $Crpt \leftarrow \emptyset$ ;  $Rvld \leftarrow \emptyset$ ;
5  $(n_C, \beta, i, b) \leftarrow \mathcal{A}^{C,R}(L)$ ;
6 if  $\exists r : b = \text{Vote}(sk_i, PK_{\mathcal{T}}, n_C, \beta, k; r) \wedge b \neq \perp \wedge b \notin Rvld \wedge sk_i \notin Crpt$  then
7   return 1
8 else
9   return 0

```

In line 1, \mathcal{A} chooses the tallier’s public key and the number of voters. Line 2 registers voters. \mathcal{A} is not permitted to influence registration while it is in progress. In particular, \mathcal{A} is not permitted to choose credential pairs, because by doing so \mathcal{A} could trivially win the experiment.

Line 4 initializes two sets: $Crpt$ is a set of voters who have been corrupted, meaning that \mathcal{A} has learned their private credential, and $Rvld$ is a set of ballots that have been revealed to \mathcal{A} . The former set models \mathcal{A} coercing voters to reveal their private credentials. The latter set models \mathcal{A} observing ballots on the bulletin board.

Line 5 challenges \mathcal{A} to produce a ballot b with the help of two oracles. Oracle C is the same oracle as in Exp-IV-Int (cf. §IV-B1); it leaks the private credentials of corrupted voters to \mathcal{A} . Oracle R reveals ballots. On invocation $R(i, \beta, n_C)$, where $1 \leq i \leq n_V$, oracle R does the following:

- Computes a ballot b that represents a vote for candidate β by a voter with private credential sk_i , that is, computes $b \leftarrow \text{Vote}(sk_i, PK_{\mathcal{T}}, n_C, \beta, k)$.
- Records b as being revealed by updating $Rvld$ to be $Rvld \cup \{b\}$.
- Outputs b .

In line 6, \mathcal{A} wins if (i) the ballot is *authentic*, meaning that it is the output of *Vote* on an authorized credential, and (ii) that credential belongs to a voter that \mathcal{A} did not corrupt, and (iii) that ballot was not revealed. If \mathcal{A} cannot succeed in this experiment, then only authorized votes are tallied.

4) *Election verifiability*: With Exp-IV-Int, Exp-UV-Int, and Exp-EV-Int, we define election verifiability with internal authentication.

Definition 9 (Ver-Int). *An election scheme Π satisfies election verifiability with internal authentication (Ver-Int) if for all PPT adversaries \mathcal{A} , there exists a negligible function μ , such that for all security parameters k , it holds that $\text{Succ}(\text{Exp-IV-Int}(\Pi, \mathcal{A}, k)) + \text{Succ}(\text{Exp-UV-Int}(\Pi, \mathcal{A}, k)) + \text{Succ}(\text{Exp-EV-Int}(\Pi, \mathcal{A}, k)) \leq \mu(k)$.*

An election scheme satisfies eligibility verifiability if $\text{Succ}(\text{Exp-EV-Int}(\Pi, \mathcal{A}, k)) \leq \mu(k)$, and similarly for individual and universal verifiability.

C. Example—Toy schemes from digital signatures

A toy election scheme satisfying Ver-Int can be based on a digital signature scheme $(\text{Gen}, \text{Sign}, \text{Ver})$. Each voter

publishes their signed candidate choice on the bulletin board.

Definition 10. *Election scheme Sig is defined as follows:*

- $\text{Setup}(k)$ outputs $(\perp, \perp, p_1(k), p_2(k))$, where p_1 and p_2 may be any polynomial functions.
- $\text{Register}(PK_{\mathcal{T}}, k)$ computes $(pk, sk) \leftarrow \text{Gen}(1^k)$ and outputs (pk, sk) .
- $\text{Vote}(sk, PK_{\mathcal{T}}, n_C, \beta, k)$ outputs $(\beta, \text{Sign}(sk, \beta))$.
- $\text{Tally}(PK_{\mathcal{T}}, SK_{\mathcal{T}}, BB, L, n_C, k)$ computes a vector \mathbf{X} of length n_C , such that \mathbf{X} is a tally of all the ballots on BB that are signed by distinct private keys whose corresponding public keys appear in L , and outputs (\mathbf{X}, \perp) .
- $\text{Verify}(PK_{\mathcal{T}}, BB, L, n_C, \mathbf{X}, P, k)$ outputs 1 if $(\mathbf{X}, P) = \text{Tally}(\perp, \perp, BB, L, n_C, \perp)$ and 0 otherwise.

The verifiability of Sig follows from the security of the underlying signature scheme:

Proposition 3. *If $(\text{Gen}, \text{Sign}, \text{Ver})$ is a signature scheme satisfying existential unforgeability under adaptive chosen-message attack, then Sig satisfies Ver-Int.*

Proof sketch. Sig satisfies individual verifiability, because voters can verify that their signed choices appear on the bulletin board. Sig satisfies universal verifiability, because signed plaintext choices are posted on BB . Finally, Sig satisfies eligibility verifiability, because anyone can check that the signed choices belong to registered voters. \square

D. Orthogonality

Exp-IV-Int, Exp-UV-Int, and Exp-EV-Int capture mostly orthogonal security properties, as shown in Table I. Individual and universal verifiability are orthogonal, and eligibility verifiability implies individual verifiability.

Theorem 4. *If an election scheme Π satisfies Exp-EV-Int, then Π also satisfies Exp-IV-Int.*

Proof sketch. If Π satisfies Exp-EV-Int, then no one can construct a ballot that appears to be associated with public credential pk unless they know private credential sk . That means that a voter can uniquely identify their ballot, because no one else knows their private credential. Therefore Π satisfies Exp-IV-Int. \square

The proof of Theorem 4 appears in the companion technical report [1].

In Table I, $\text{AlwaysVerify}(\cdot)$ is a function that transforms an election scheme by compromising Verify to always return 1. Thus, $\text{AlwaysVerify}(\Pi)$ is guaranteed not to satisfy Exp-UV-Int. Similarly, $\text{IgnoreCreds}(\cdot)$ is a function that accepts as input an election scheme with external authentication and returns as output an election scheme with internal authentication. The resulting scheme, however, simply ignores credentials altogether: Register returns (\perp, \perp) , Vote ignores sk , and Tally and Verify ignore L . Thus, $\text{IgnoreCreds}(\Pi)$ is guaranteed not to satisfy Exp-EV-Int. Using those functions, we briefly explain each line of the table:

Line	IV	UV	EV	Scheme
1	\times	\times	\times	$\text{AlwaysVerify}(\text{IgnoreCreds}(\text{Choice}))$
2	\times	\times	\checkmark	—
3	\times	\checkmark	\times	$\text{IgnoreCreds}(\text{Choice})$
4	\times	\checkmark	\checkmark	—
5	\checkmark	\times	\times	$\text{AlwaysVerify}(\text{IgnoreCreds}(\text{Nonce}))$
6	\checkmark	\times	\checkmark	$\text{AlwaysVerify}(\text{Sig})$
7	\checkmark	\checkmark	\times	Malleable Sig
8	\checkmark	\checkmark	\checkmark	Sig

TABLE I

ELECTION SCHEMES THAT SATISFY EACH COMBINATION OF INDIVIDUAL, UNIVERSAL AND ELIGIBILITY VERIFIABILITY

- 1) Recall (from §II-D) that Choice is the election scheme in which ballots contain only the plaintext candidate choice. By compromising Verify and ignoring credentials, we obtain a scheme that satisfies no properties.
- 2) By Theorem 4, this situation is impossible.
- 3) Compared to line 1 of Table I, this scheme satisfies Exp-UV-Int, because Verify is not compromised.
- 4) By Theorem 4, this situation is impossible.
- 5) Nonce satisfies Exp-IV-Ext and Exp-UV-Ext. Moreover, $\text{IgnoreCreds}(\text{Nonce})$ satisfies Exp-IV-Int and Exp-UV-Int. By compromising Verify, we obtain a scheme that satisfies only Exp-IV-Int.
- 6) Sig satisfies all three properties. By compromising Verify, we obtain a scheme that satisfies only Exp-IV-Int and Exp-EV-Int.
- 7) By making Sig 's underlying signature scheme malleable,²⁰ we could obtain a scheme that does not satisfy Exp-EV-Int, because the adversary could construct a valid ballot out of a revealed ballot. But the scheme would continue to satisfy Exp-IV-Int and Exp-UV-Int.
- 8) Sig satisfies all three properties.

V. CASE STUDY: JCJ

JCJ (named for its designers, Juels, Catalano, and Jakobsen) [57]–[59] is a *coercion-resistant* election scheme, meaning voters cannot prove whether or how they voted, even if they can interact with the adversary while voting. Coercion resistance protects elections from improper influence by adversaries.

To achieve verifiability and coercion resistance, JCJ uses verifiable *mixnets*, which anonymize a set of messages.²¹ During tallying, all encrypted choices are anonymized by a mixnet, then all choices are decrypted. The tally is computed from the decrypted choices. Informally, JCJ works as follows:

- **Setup.** The tallier generates a key pair $(PK_{\mathcal{T}}, SK_{\mathcal{T}})$ for an encryption scheme and publishes the public key.
- **Registration.** To register a voter, the registrar generates a nonce, which is sent to the voter and serves as the private credential. The public credential is computed as an

²⁰Given a message m and signature σ , a *malleable* signature scheme permits computation of a signature σ' on a related message m' [21]. The malleable signature scheme Sig used in line 7 of Table I would need to enable an adversary to transform a signature on a well-formed candidate β into a signature on a distinct, well-formed candidate β' .

²¹Chaum [22] introduced mixnets. Adida [2] surveys verifiable mixnets.

encryption of the private credential with $PK_{\mathcal{T}}$. After all voters are registered, the registrar publishes the electoral roll.

- **Voting.** A voter encrypts her candidate choice with $PK_{\mathcal{T}}$. She also encrypts her private credential with $PK_{\mathcal{T}}$. She proves in zero-knowledge that she simultaneously knows both plaintexts, and that her choice is well-formed. The voter posts her ballot (i.e., both ciphertexts and the proof) on the bulletin board.
- **Tallying.** The tallier discards any ballots from the bulletin board for which the zero-knowledge proofs do not verify. All unauthorized ballots are then discarded through a combination of protocols that includes verifiable mixnets and *plaintext equivalence tests* (PETs) [54]. (PETs enable proof that two ciphertexts contain the same plaintext without revealing that plaintext.) The tallier decrypts and publishes the remaining ballots, along with a proof that decryption was performed correctly.
- **Verification.** A verifier checks all the proofs included in ballots, and all the proofs published during tallying.

The companion technical report [1] gives a formal description of JCJ. That formalization satisfies individual and universal verifiability, assuming that the cryptographic primitives satisfy certain properties that we identify. But the formalization fails to satisfy eligibility verifiability, because knowledge of the tallier’s private key $SK_{\mathcal{T}}$ suffices to construct ballots that appear authentic: with $SK_{\mathcal{T}}$, any public credential can be decrypted to discover the corresponding private credential. Note that Exp-EV-Int permits an adversary \mathcal{A} to choose the tallier’s key pair, so \mathcal{A} does know $SK_{\mathcal{T}}$ hence can construct a ballot that suffices to win Exp-EV-Int.

We can nonetheless prove that JCJ satisfies a variant of eligibility verifiability. Consider the following experiment, which does not permit the adversary to choose the tallier’s key pair:

```

Exp-EV-Int-Weak( $\Pi, \mathcal{A}, k$ ) =
1  $(PK_{\mathcal{T}}, SK_{\mathcal{T}}, m_B, m_C) \leftarrow \text{Setup}(k);$ 
2  $n_V \leftarrow \mathcal{A}(PK_{\mathcal{T}}, k);$ 
3 for  $1 \leq i \leq n_V$  do  $(pk_i, sk_i) \leftarrow \text{Register}(PK_{\mathcal{T}}, k);$ 
4  $L \leftarrow \{pk_1, \dots, pk_{n_V}\};$ 
5  $Crpt \leftarrow \emptyset; Rvld \leftarrow \emptyset;$ 
6  $(n_C, \beta, i, b) \leftarrow \mathcal{A}^{C,R}(L);$ 
7 if  $\exists r : b = \text{Vote}(sk_i, PK_{\mathcal{T}}, n_C, \beta, k; r) \wedge b \neq \perp \wedge b \notin Rvld \wedge sk_i \notin Crpt$  then
8   return 1
9 else
10  return 0

```

Line 1 of Exp-EV-Int has been refactored into lines 1 and 2 of Exp-EV-Int-Weak. In line 1 of Exp-EV-Int-Weak, keys are generated by the experiment. In line 2, \mathcal{A} is given the public key but not the private key.

Using Exp-EV-Int-Weak, we define a weaker variant of Ver-Int and prove that JCJ satisfies it:

Definition 11 (Ver-Int-Weak). *An election scheme Π sat-*

isfies weak election verifiability with internal authentication (Ver-Int-Weak) if for all probabilistic polynomial-time adversaries \mathcal{A} , there exists a negligible function μ , such that for all security parameters k , we have $\text{Succ}(\text{Exp-IV-Int}(\Pi, \mathcal{A}, k)) + \text{Succ}(\text{Exp-UV-Int}(\Pi, \mathcal{A}, k)) + \text{Succ}(\text{Exp-EV-Int-Weak}(\Pi, \mathcal{A}, k)) \leq \mu(k)$.

Theorem 5. *JCJ satisfies Ver-Int-Weak.*

Proof sketch. JCJ satisfies individual verifiability, because the probabilistic encryption scheme ensures that ballots are unique, with overwhelming probability. JCJ satisfies universal verifiability, because the proofs produced throughout tallying can be publicly verified. And JCJ satisfies eligibility verifiability, because \mathcal{A} cannot construct new ballots without knowing a voter’s private credential or the tallier’s private key. \square

A formal proof of Theorem 5 appears in the companion technical report [1]. The proof assumes the random oracle model.

The Civitas [29] scheme refines the JCJ scheme. Some refinements relevant to election verifiability are an implementation of a distributed registration protocol, and a mixnet based on randomized partial checking (RPC) [55]. We leave a proof that Civitas satisfies Ver-Int-Weak as future work.

VI. NEW CLASSES OF ATTACK

Our definitions of election verifiability improve upon existing definitions by detecting two previously unidentified classes of attack:

- **Collusion attacks.** An election scheme’s tallying and verification algorithms might be designed such that they collude to accept incorrect tallies.
- **Biassing attacks.** An election scheme’s verification algorithm might be designed such that it rejects some legitimate tallies.

Although a well-designed election scheme would hopefully not exhibit these vulnerabilities, it is the job of verifiability definitions to detect malicious schemes, regardless of whether vulnerabilities are due to malice or errors. So definitions of election verifiability should preclude collusion and biassing attacks.

A. Collusion Attacks

Here are two examples of potential collusion attacks:

- **Vote stuffing.** Tally behaves normally, but adds κ votes for candidate β . Verify subtracts κ votes from β , then proceeds with verification as normal. Elections thus verify as normal, except that candidate β receives extra votes.
- **Backdoor tally replacement.** Tally and Verify behave normally, unless a *backdoor* value is posted on the bulletin board BB . For example, if $(SK_{\mathcal{T}}, X^*)$ appears on BB , then Tally and Verify both ignore the correct tally and instead replace it with tally X^* . Value $SK_{\mathcal{T}}$ is the backdoor here; it cannot appear on BB (except with negligible probability) unless the tallier is malicious.

Vote stuffing is detected by our definitions of Correctness (§II-A and §IV-A), because these definitions require that the tally produced by Tally corresponds to the choices encapsulated in ballots on the bulletin board. Note that vote stuffing is not a failure of eligibility verifiability, because the stuffed votes do not correspond to any ballots on the bulletin board. Backdoor tally replacement is detected by our definitions of universal verifiability (§II-B2 and §IV-B2), because those definitions require Verify to accept only those tallies that correspond to a correct tally of the bulletin board.

We show, next, that the definition of election verifiability by Juels et al. [59] fails to detect vote stuffing and backdoor tally replacement, and that the definition by Cortier et al. [32] fails to detect backdoor tally replacement.

Juels et al. [59] formalize definitions that we name *JCJ-correctness* and *JCJ-verifiability*. JCJ-correctness is intuitively meant to capture that “ \mathcal{A} cannot pre-empt, alter, or cancel the votes of honest voters [and] that \mathcal{A} cannot cause voters to cast ballots resulting in double voting” [59, p. 45]; it is formalized in terms of whether the adversary can post ballots on the bulletin board that cause the tally to be computed incorrectly. JCJ-verifiability is intuitively “the ability for any player to check whether the tally...has been correctly computed” [59, p. 46]; it is formalized in terms of whether Verify will accept a tally that differs from the output of Tally. We restate the formal definitions in the companion technical report [1].

To show that the JCJ definitions fail to detect collusion attacks, we first formalize the vote stuffing attack. An election scheme $\Pi = (\dots, \text{Tally}, \text{Verify})$ can be modified to derive a vote-stuffing election scheme $\text{Stuff}(\Pi, \beta, \kappa) = (\dots, \text{Tally}_S, \text{Verify}_S)$, which adds κ votes to β , as follows:²²

- $\text{Tally}_S(\dots)$ computes $(\mathbf{X}, P) \leftarrow \text{Tally}(\dots)$ and outputs $(\text{Add}(\mathbf{X}, \beta, \kappa), P)$.²³
- $\text{Verify}_S(\dots, \mathbf{X}, P, \dots)$ computes $v \leftarrow \text{Verify}(\dots, \text{Add}(\mathbf{X}, \beta, -\kappa), P, \dots)$ and outputs v .

Proposition 6. *Given an election scheme Π satisfying JCJ-correctness and JCJ-verifiability, and integers β and κ , it holds that $\text{Stuff}(\Pi, \beta, \kappa)$ satisfies JCJ-correctness and JCJ-verifiability.*

A formal proof appears in the companion technical report [1].

Second, we formalize the backdoor tally replacement attack. An election scheme $\Pi = (\dots, \text{Tally}, \text{Verify})$ can be modified to derive an election scheme $\text{Backdoor}(\Pi) = (\dots, \text{Tally}_B, \text{Verify}_B)$, as follows:

- $\text{Tally}_B(\dots, SK_{\mathcal{T}}, BB, \dots)$ checks whether there exists a unique \mathbf{X}^* such that $(SK_{\mathcal{T}}, \mathbf{X}^*) \in BB$. If so, Tally_B outputs (\mathbf{X}^*, \perp) . Otherwise, Tally_B computes $(\mathbf{X}, P) \leftarrow \text{Tally}(\dots, SK_{\mathcal{T}}, BB, \dots)$ and outputs (\mathbf{X}, P) .
- $\text{Verify}_B(PK_{\mathcal{T}}, \dots, BB, \dots, \mathbf{X}, P, \dots)$ checks whether

²²We omit many of the parameters of Tally and Verify here for simplicity; see the companion technical report [1] for details.

²³Let $\text{Add}(\mathbf{X}, \beta, \kappa) = (\mathbf{X}[1], \dots, \mathbf{X}[\beta - 1], \mathbf{X}[\beta] + \kappa, \mathbf{X}[\beta + 1], \dots, \mathbf{X}[|\mathbf{X}|])$. And let $|\mathbf{X}|$ denote the length of vector \mathbf{X} .

there exists a unique \mathbf{X}^* such that $(SK_{\mathcal{T}}, \mathbf{X}^*) \in BB$.²⁴ If so, Verify_B outputs 1. Otherwise, Verify_B computes $v \leftarrow \text{Verify}(PK_{\mathcal{T}}, \dots, BB, \dots, \mathbf{X}, P, \dots)$ and outputs v .

Proposition 7. *Given an election scheme Π satisfying JCJ-correctness and JCJ-verifiability that does not leak the tallyer’s private key, it holds that $\text{Backdoor}(\Pi)$ satisfies JCJ-correctness and JCJ-verifiability.*

A formal proof appears in the companion technical report [1], where we also formally define key leakage.

Cortier et al. [32] propose definitions similar to *JCJ-verifiability* and insist that election schemes must satisfy their notions of correctness and partial tallying. Vote stuffing is detected by their correctness property, but backdoor tally replacement is not. The ideas remain the same, so we omit formalized results. We have reported these findings to the original authors [31], [43], [44].

B. Biasing attacks

Here are three formalizations of biasing attacks, derived from an election scheme $\Pi = (\dots, \text{Verify})$.

- **Reject All.** Let $\text{Reject}(\Pi)$ be (\dots, Verify_R) , where Verify_R always outputs 0. Verify_R therefore always rejects, hence no election can ever be considered valid.
- **Selective Reject.** Let ε be a distinguished value that would not be posted on the bulletin board by honest voters. Let $\text{Selective}(\Pi, \varepsilon)$ be (\dots, Verify_R) , where $\text{Verify}_R(\dots, BB, \dots)$ computes $v \leftarrow \text{Verify}(\dots, BB, \dots)$ and outputs 1 if both $v = 1$ and $\varepsilon \notin BB$. Otherwise, Verify_R outputs 0. Verify_R therefore rejects if ε appears on the bulletin board, hence some elections can be invalidated.
- **Biased Reject.** Suppose Z is a set of tallies. Let $\text{Bias}(\Pi, Z)$ be (\dots, Verify_R) , where $\text{Verify}_R(\dots, \mathbf{X}, \dots)$ computes $v \leftarrow \text{Verify}(\dots, \mathbf{X}, \dots)$ and outputs 1 if both $v = 1$ and $\mathbf{X} \in Z$. Otherwise, Verify_R outputs 0. Verify_R therefore only accepts a subset of the tallies accepted by Verify, hence biases tallies toward Z .

These formalizations do not satisfy our definition of Completeness (§II-A and §IV-A), hence, our definitions of verifiability detect these biasing attacks.

The definition of verifiability by Juels et al. [59] fails to detect all three of the above attacks, because that definition has no notion of Completeness. For example, it is vulnerable to Biased Reject attacks:

Proposition 8. *Given an election scheme Π satisfying JCJ-correctness and JCJ-verifiability, and given a multiset Z , it holds that $\text{Bias}(\Pi, Z)$ satisfies JCJ-correctness and JCJ-verifiability.*

A formal proof appears in the companion technical report [1].

The definition of verifiability by Kiayias et al. [61] fails to detect Selective Reject attacks, because (like JCJ) the

²⁴ Verify_B also needs to check that $SK_{\mathcal{T}}$ is the private key corresponding to $PK_{\mathcal{T}}$. We omit formalizing this detail, but note that it is straightforward for real-world encryption schemes such as El Gamal and RSA.

definition has no notion of Completeness. Their notion of Correctness does rule out Reject All and Biased Reject attacks.

Similarly, the definition of verifiability by Cortier et al. [32] detects Biased Reject and Reject All attacks, but fails to detect Selective Reject attacks, because that definition’s notion of Completeness does not quantify over all bulletin boards.

VII. RELATED WORK

Kiayias [60] presents an overview of security properties for election schemes. Many election schemes in the literature state properties called correctness, accuracy, or (universal) verifiability without formally defining those terms.

In the computational model, Juels et al. [57]–[59] and Cortier et al. [32] give game-based definitions of verifiability. Those definitions fail to detect biasing and collusion attacks (cf. §VI). Definitions of universal verifiability (which is just one aspect of election verifiability) in the computational model seem to originate with Benaloh and Tuinstra [13], who define a *correctness* property that says every participant is convinced that the tally is accurate with respect to the votes cast, and with Cohen and Fischer [30], who define *verifiability* to mean that there exists a *check* function that returns *good* iff the announced tally of the election corresponds to the cast votes.

Kiayias et al. [61] define a property they name *E2E verifiability* (E2E abbreviates “end-to-end”). This property combines our intuitive notions of individual and universal verifiability into a single definition. Their definition fails to detect Selective Reject attacks (cf. §VI). Their definitions, like ours, do not address voter intent—that is, verification by humans that ballots correctly encode candidate choices—as we discuss in Section VIII.

Also in the computational model, Groth [46], and Moran and Naor [72], state definitions of verifiability in terms of *universal composability* [20]. These definitions involve defining an *ideal functionality*; part of that is similar to our *correct-tally* function. Groth’s definition does not guarantee universal verifiability [46, p. 2], but Moran and Naor’s does [72, p. 386].

In the symbolic model, Smyth et al. [83] define the first definition of election verifiability. This definition is amenable to automated reasoning, but is stronger than necessary and cannot be satisfied by many election schemes, including Helios and Civitas. Kremer et al. [63] overcome this limitation with a weaker definition that sacrifices amenability to automated reasoning, and Smyth [80, §3] extends this definition. Dreier et al. have adapted election verifiability to auction [40] and examination [39] systems.

Also in the symbolic model, Kremer and Ryan [62] and Backes et al. [8] formalize definitions of *eligibility*. These definitions are not intended to provide assurances if the election authorities are dishonest. For example, the definition of Kremer and Ryan does not detect whether corrupt election authorities insert votes [62, §5.2]. Likewise, the definition of Backes et al. assumes that election authorities are honest [8, §3].

Our definition of election verifiability follows Smyth et al. [63], [80], [83] by deconstructing it into individual, universal, and eligibility verifiability. Other deconstructions of election verifiability are possible. For example, Adida and Neff [7, §2] identify four aspects of verifiability:

- *Cast as intended*: the ballot is cast at the polling station as the voter intended.
- *Recorded as cast*: cast ballots are preserved with integrity through the ballot collection process.
- *Counted as recorded*: recorded ballots are counted correctly.
- *Eligible voter verification*: only eligible voters can cast a ballot in the first place.

Those definitions are not mathematical, so we cannot attempt a precise comparison. Nonetheless, eligibility verifiability and eligible voter verification seem to be addressing similar concerns. Likewise, individual and universal verifiability together seem to be addressing concerns similar to that of recorded as cast and counted as recorded together. Recorded as cast, in our work, reduces to the bulletin board preserving ballots with integrity—a property that we have assumed, because cryptographic election schemes assume it, too. Ways to construct secure bulletin boards have been proposed, e.g., [36], [48], [76], [78]. We postpone a discussion of cast as intended to Section VIII.

Privacy properties [38], [59], [69], [70], [81], [82]—such as ballot secrecy, receipt freeness, and coercion resistance—complement verifiability. Chevallier-Mames et al. [26], [27] and Hosp and Vora [51], [52] show an incompatibility result: election schemes cannot unconditionally satisfy privacy and universal verifiability. But weaker versions of these properties can hold simultaneously, as can be witnessed from Theorems 2 and 5 coupled with existing privacy results such as the ballot secrecy proofs for Helios 4.0 [18, Theorem 3], [16, Theorem 6.12], and the coercion resistance proof for JCJ [59, §5].

Comparison with global verifiability: Küsters et al. [68], [69], [71] present a definition of *global verifiability* that can be used with any kind of protocol, not just electronic voting protocols. To analyze the verifiability of a protocol, users of this definition must themselves formalize *goals*, which are properties required to hold in every run of the protocol. For example, a goal γ_e is presented in a case study [69, §5.2] of global verifiability applied to voting:

γ_e contains all runs for which there exist choices of the dishonest voters (where a choice is either to abstain or to vote for one of the candidates) such that the result obtained together with the choices made by the honest voters in this run differs only by ℓ votes from the published result (i.e. the result that can be computed from the simple ballots on the bulletin board).

Another goal γ is presented in a case study [71, §6.2] of Helios:

γ is satisfied in a run if the published result exactly

reflects the actual votes of the honest voters in this run and votes of dishonest voters are distributed in some way on the candidates, possibly in a different way than how the dishonest voters actually voted.

These informal statements of goals are appealing, but they do not constitute rigorous mathematical definitions. As Kiayias et al. write, “[global verifiability] has the disadvantage that the set γ remains undetermined and thus the level of verifiability that is offered by the definition hinges on the proper definition of γ which may not be simple” [61, p. 476]. In our own work, we found that formal definitions were quite tricky to get right—for example, which ballots should be counted, how to count them, and how to determine whether that count differed from the published tally. So we shared [65] and discussed [66] our results with Küsters. In response, Küsters et al. updated an online technical report to propose a formalization of goals [64, §5.2]; we look forward to analyzing that formalization when it is published.

In an analysis of Helios, Küsters et al. [71] use goal γ to conclude that Helios 2.0 satisfies global verifiability. Yet Bernhard et al. [18] demonstrate an attack against the verifiability of Helios 2.0, and in the companion technical report [1] we show that Helios 2.0 does not satisfy Ver-Ext. This seeming discrepancy arises because the analysis in [71] does not formalize all the cryptographic primitives used by Helios, hence the attack goes unnoticed. So another contribution of our own work is to correctly distinguish between unverifiable and verifiable variants of Helios by rigorously analyzing the cryptography used in Helios.

It is natural to ask whether election verifiability can be expressed in terms of global verifiability. We believe it can be. For instance, individual, universal and eligibility verifiability could be expressed, in the informal style of the goals quoted above, as the following goals:

- γ_{IV} is satisfied in a run if voters can uniquely identify their ballots on the bulletin board in this run.
- γ_{UV} is satisfied in a run if the correct tally of votes cast by authorized voters in this run is the same as the tally produced by algorithm Tally.
- γ_{EV} is satisfied in a run if every ballot tallied in this run was created by a voter in possession of a private credential.

Küsters et al. [69] argue that deconstructing verifiability into individual and universal verifiability is insufficient to detect certain attacks involving ill-formed ballots. But those attacks leave open the possibility that there do exist notions of individual and universal verifiability that would be sufficient. Indeed, our own definition of universal verifiability rules out attacks based on ill-formed ballots, because *correct-tally* ensures that tallied ballots are well-formed.

One concern that might be raised is whether there still lurk any “gaps” in our decomposition into individual and universal (and eligibility) verifiability. Indeed, there might be. But the definition of global verifiability does not rule out the possibility of gaps, either: any gap in the formal statement of a

goal will lead to a vulnerability. That is, if the analyst forgets to include some necessary facet of verifiability when stating the formal goal, then global verifiability will not detect any attacks against that facet. Global verifiability does not guarantee a lack of gaps.

VIII. CONCLUDING REMARKS

When we began this work, we were studying the Juels et al. [59] definition of election verifiability. We discovered that the definition fails to detect biasing and collusion attacks. While attempting to improve the Juels et al. definition to rule out those attacks, we discovered that factoring it into individual, universal, and eligibility verifiability led to an elegant decomposition of (mostly) orthogonal properties. We later sought to apply our new definitions to existing electronic voting systems, and Helios [6] and Civitas [29] were natural choices. But they treat authentication differently—Helios outsources authentication, whereas Civitas does not—so we were led to separate our definitions into variants for external and internal authentication. We were at first surprised to discover that JCJ, hence Civitas, does not satisfy the strong definition of eligibility verifiability. But upon reflection, it became apparent that an adversary who knows the taller’s private key can easily forge ballots that appear to be from eligible voters.

Our definitions of verifiability have not addressed the issue of voter intent—that is, verification by a human that the ballot submitted by a voter corresponds to the candidate choice the voter intended to make. Adida and Neff call this property “cast as intended” [7]. Many election schemes (e.g., [42], [50], [59], [61]) do not satisfy cast as intended, because the schemes implicitly or explicitly assume that voters can themselves verify the cryptographic operations required to construct ballots. Nevertheless, schemes by Chaum [23], Neff [73], and Benaloh [11], [12] introduce cryptographic mechanisms to verify voter intent. It would be natural to explore strengthening our definitions to address voter intent.

The goal of this research is to enable verifiability of the voting systems we use in real-life, rather than merely trusting them. Research on verifiability can generalize beyond voting to other systems that must guarantee strong forms of integrity. Verifiable voting systems thus have the potential to contribute to the science of security, to democracy, and to broader society.

ACKNOWLEDGMENTS

We thank David Bernhard, Jeremy Clark, Véronique Cortier, David Galindo, Markus Jakobsson, Steve Kremer, Ralf Küsters, Elizabeth Quaglia, Mark Ryan, Susan Thomson, and Poorvi Vora for insightful discussions that have influenced this paper. This work is partly supported by the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013) / ERC project CRYSP (259639), by AFOSR grants FA9550-12-1-0334 and FA9550-14-1-0334, by NSF grant 1421373, and by the National Security Agency. This work was performed in part at George Washington University and INRIA.

DEDICATION

Ben Smyth dedicates his contribution to the loving memory of Anne Konishi, 1971 – 2015. What matters most of all is the dash. We had a great time.

He writes for Christina Mai Konishi. Smile like your mother, for good fortune seeks those who smile (*warau kado niwa fuku kitaru*, says the Japanese proverb).

REFERENCES

- [1] Election verifiability: Cryptographic definitions and an analysis of Helios and JCJ (technical report), August 2015. Available from https://drive.google.com/file/d/0B_LA8FeLbZ4xMEdNeThRVS1KYTA/view.
- [2] Ben Adida. *Advances in Cryptographic Voting Systems*. PhD thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, 2006.
- [3] Ben Adida. Helios: Web-based Open-Audit Voting. In *USENIX Security'08: 17th USENIX Security Symposium*, pages 335–348. USENIX Association, 2008.
- [4] Ben Adida. Helios deployed at Princeton. <http://heliosvoting.wordpress.com/2009/10/13/helios-deployed-at-princeton/> (accessed 7 May 2014), 2009.
- [5] Ben Adida. Helios v4 Verification Specs. Helios documentation, <http://documentation.heliosvoting.org/verification-specs/helios-v4> (accessed 2 May 2014), 2014. A snapshot of the specification on 8 Oct 2013 is available from <https://web.archive.org/web/20131018033747/http://documentation.heliosvoting.org/verification-specs/helios-v4>.
- [6] Ben Adida, Olivier de Marneffe, Olivier Pereira, and Jean-Jacques Quisquater. Electing a University President Using Open-Audit Voting: Analysis of Real-World Use of Helios. In *EVT/WOTE'09: Electronic Voting Technology Workshop/Workshop on Trustworthy Elections*. USENIX Association, 2009.
- [7] Ben Adida and C. Andrew Neff. Ballot casting assurance. In *EVT'06: Electronic Voting Technology Workshop*, Berkeley, CA, USA, 2006. USENIX Association.
- [8] Michael Backes, Cătălin Hrițcu, and Matteo Maffei. Automated Verification of Remote Electronic Voting Protocols in the Applied Pi-calculus. In *CSF'08: 21st Computer Security Foundations Symposium*, pages 195–209. IEEE Computer Society, 2008.
- [9] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *Proceedings of the 1st ACM Conference on Computer and Communications Security*, CCS '93, pages 62–73, New York, NY, USA, 1993. ACM.
- [10] Josh Benaloh. *Verifiable Secret-Ballot Elections*. PhD thesis, Department of Computer Science, Yale University, 1996.
- [11] Josh Benaloh. Simple Verifiable Elections. In *EVT'06: Electronic Voting Technology Workshop*. USENIX Association, 2006.
- [12] Josh Benaloh. Ballot Casting Assurance via Voter-Initiated Poll Station Auditing. In *EVT'07: Electronic Voting Technology Workshop*. USENIX Association, 2007.
- [13] Josh Benaloh and Dwight Tuinstra. Receipt-free secret-ballot elections (extended abstract). In *Proceedings of the Twenty-sixth Annual ACM Symposium on Theory of Computing*, STOC '94, pages 544–553, New York, NY, USA, 1994. ACM.
- [14] Josh Benaloh, Serge Vaudenay, and Jean-Jacques Quisquater. Final Report of IACR Electronic Voting Committee. International Association for Cryptologic Research. http://www.iacr.org/elections/eVoting/finalReportHelios_2010-09-27.html (accessed 7 May 2014), Sept 2010.
- [15] Josh Benaloh and Moti Yung. Distributing the Power of a Government to Enhance the Privacy of Voters. In *PODC'86: 5th Principles of Distributed Computing Symposium*, pages 52–62. ACM Press, 1986.
- [16] David Bernhard. *Zero-Knowledge Proofs in Theory and Practice*. PhD thesis, Department of Computer Science, University of Bristol, 2014.
- [17] David Bernhard, Véronique Cortier, Olivier Pereira, Ben Smyth, and Bogdan Warinschi. Adapting Helios for provable ballot privacy. In *ESORICS'11: 16th European Symposium on Research in Computer Security*, volume 6879 of *LNCS*, pages 335–354. Springer, 2011.
- [18] David Bernhard, Olivier Pereira, and Bogdan Warinschi. How Not to Prove Yourself: Pitfalls of the Fiat-Shamir Heuristic and Applications to Helios. In *ASIACRYPT'12: 18th International Conference on the Theory and Application of Cryptology and Information Security*, volume 7658 of *LNCS*, pages 626–643. Springer, 2012.
- [19] Debra Bowen. Secretary of State Debra Bowen Moves to Strengthen Voter Confidence in Election Security Following Top-to-Bottom Review of Voting Systems. California Secretary of State, press release DB07:042 <http://www.sos.ca.gov/voting-systems/oversight/tbr/db07-042-tbbr-system-decisions-release.pdf> (accessed 7 May 2014), August 2007.
- [20] Ran Canetti. Universally composable security: a new paradigm for cryptographic protocols. In *42nd IEEE Symposium on Foundations of Computer Science*, pages 136–145, October 2001.
- [21] Melissa Chase, Markulf Kohlweiss, Anna Lysyanskaya, and Sarah Meiklejohn. Malleable signatures: New definitions and delegatable anonymous credentials. In *CSF'14: 27th Computer Security Foundations Symposium*. IEEE Computer Society, 2014. To appear.
- [22] David Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–88, 1981.
- [23] David Chaum. Secret-ballot receipts: True voter-verifiable elections. *IEEE Security and Privacy*, 2(1):38–47, 2004.
- [24] David Chaum, Richard Carback, Jeremy Clark, Aleksander Essex, Stefan Popoveniuc, Ronald L. Rivest, Peter Y. A. Ryan, Emily Shen, and Alan T. Sherman. Scantegrity II: End-to-end verifiability for optical scan election systems using invisible ink confirmation codes. In *Proc. Conference on Electronic Voting Technology*, pages 14:1–14:13. USENIX, 2008.
- [25] David Chaum, Peter Y. A. Ryan, and Steve Schneider. A Practical Voter-Verifiable Election Scheme. In *ESORICS'05: 10th European Symposium On Research In Computer Security*, volume 3679 of *LNCS*, pages 118–139. Springer, 2005.
- [26] Benoît Chevallier-Mames, Pierre-Alain Fouque, David Pointcheval, Julien Stern, and Jacques Traoré. On Some Incompatible Properties of Voting Schemes. In *WOTE'06: Workshop on Trustworthy Elections*, 2006.
- [27] Benoît Chevallier-Mames, Pierre-Alain Fouque, David Pointcheval, Julien Stern, and Jacques Traoré. On Some Incompatible Properties of Voting Schemes. In David Chaum, Markus Jakobsson, Ronald L. Rivest, and Peter Y. A. Ryan, editors, *Towards Trustworthy Elections: New Directions in Electronic Voting*, volume 6000 of *LNCS*, pages 191–199. Springer, 2010.
- [28] Michael Clarkson, Brian Hay, Meador Inge, abhi shelat, David Wagner, and Alec Yasinsac. Software review and security analysis of Scyll remote voting software. Report commissioned by Florida Division of Elections. Available from <http://election.dos.state.fl.us/voting-systems/pdf/FinalReportSept19.pdf>. Filed September 19, 2008.
- [29] Michael R. Clarkson, Stephen Chong, and Andrew C. Myers. Civitas: Toward a Secure Voting System. In *S&P'08: 29th Security and Privacy Symposium*, pages 354–368. IEEE Computer Society, 2008.
- [30] Josh Daniel Cohen and Michael J. Fischer. A Robust and Verifiable Cryptographically Secure Election Scheme. In *FOCS'85: 26th Symposium on Foundations of Computer Science*, pages 372–382. IEEE Computer Society, 1985.
- [31] Véronique Cortier and David Galindo. Private communication, Nancy, France, 13th June 2013.
- [32] Véronique Cortier, David Galindo, Stephane Gloudu, and Malika Izabachène. Election Verifiability for Helios under Weaker Trust Assumptions. In *ESORICS'14: 19th European Symposium on Research in Computer Security*, volume 8713 of *LNCS*, pages 327–344. Springer, 2014.
- [33] Véronique Cortier and Ben Smyth. Attacking and fixing Helios: An analysis of ballot secrecy. In *CSF'11: 24th Computer Security Foundations Symposium*, pages 297–311. IEEE Computer Society, 2011.
- [34] Véronique Cortier and Ben Smyth. Attacking and fixing Helios: An analysis of ballot secrecy. *Journal of Computer Security*, 21(1):89–148, 2013.
- [35] Ronald Cramer, Rosario Gennaro, and Berry Schoenmakers. A Secure and Optimally Efficient Multi-Authority Election Scheme. In *EUROCRYPT'97: 16th International Conference on the Theory and Applications of Cryptographic Techniques*, volume 1233 of *LNCS*, pages 103–118. Springer, 1997.
- [36] Chris Culnane and Steve A. Schneider. A Peered Bulletin Board for Robust Use in Verifiable Voting Systems. In *CSF'14: 27th Computer*

The dedication references Linda Ellis (1996) *The Dash*.

- Security Foundations Symposium*, pages 169–183. IEEE Computer Society, 2014.
- [37] Ivan Damgård, Mads Jurik, and Jesper Buus Nielsen. A Generalization of Paillier’s Public-Key System with Applications to Electronic Voting. *International Journal of Information Security*, 9(6):371–385, 2010.
 - [38] Stéphanie Delaune, Steve Kremer, and Mark D. Ryan. Verifying privacy-type properties of electronic voting protocols. *Journal of Computer Security*, 17(4):435–487, July 2009.
 - [39] Jannik Dreier, Rosario Giustolisi, Ali Kassem, Pascal Lafourcade, and Gabriele Lenzini. On the Verifiability of (Electronic) Exams. Technical Report TR-2014-2, Verimag, 2014.
 - [40] Jannik Dreier, Hugo Jonker, and Pascal Lafourcade. Defining verifiability in e-auction protocols. In *ASIACCS’13: 8th ACM Symposium on Information, Computer and Communications Security*, pages 547–552. ACM Press, 2013.
 - [41] Taher ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Transactions on Information Theory*, 31(4):469–472, 1985.
 - [42] Atsushi Fujioka, Tatsuaki Okamoto, and Kazuo Ohta. A Practical Secret Voting Scheme for Large Scale Elections. In *AUSCRYPT’92: Workshop on the Theory and Application of Cryptographic Techniques*, volume 718 of *LNCS*, pages 244–251. Springer, 1992.
 - [43] David Galindo and Véronique Cortier. Private email communication, 19th June 2013.
 - [44] David Galindo and Véronique Cortier. Private email communication, Summer/Autumn 2014.
 - [45] Use of voting computers in 2005 Bundestag election unconstitutional, March 2009. Press release 19/2009 <http://www.bundesverfassungsgericht.de/en/press/bvg09-019en.html> (accessed 7 May 2014).
 - [46] Jens Groth. Evaluating security of voting schemes in the universal compositability framework. In *Applied Cryptography and Network Security*, volume 3089 of *Lecture Notes in Computer Science*, pages 46–60. Springer, 2004.
 - [47] Stuart Haber, Josh Benaloh, and Shai Halevi. The Helios e-Voting Demo for the IACR. International Association for Cryptologic Research. <http://www.iacr.org/elections/eVoting/heliosDemo.pdf> (accessed 7 May 2014), May 2010.
 - [48] James Heather and David Lundin. The Append-Only Web Bulletin Board. In *FAST’08: 5th International Workshop on Formal Aspects in Security and Trust*, volume 5491 of *LNCS*, pages 242–256. Springer, 2008.
 - [49] Martin Hirt. Receipt-Free K -out-of- L Voting Based on ElGamal Encryption. In David Chaum, Markus Jakobsson, Ronald L. Rivest, and Peter Y. A. Ryan, editors, *Towards Trustworthy Elections: New Directions in Electronic Voting*, volume 6000 of *LNCS*, pages 64–82. Springer, 2010.
 - [50] Martin Hirt and Kazue Sako. Efficient Receipt-Free Voting Based on Homomorphic Encryption. In *EUROCRYPT’06: 25th International Conference on the Theory and Applications of Cryptographic Techniques*, volume 1807 of *LNCS*, pages 539–556. Springer, 2006.
 - [51] Benjamin Hosp and Poorvi L. Vora. An information-theoretic model of voting systems. In *WOTE’06: Workshop on Trustworthy Elections*, 2006.
 - [52] Benjamin Hosp and Poorvi L. Vora. An information-theoretic model of voting systems. *Mathematical and Computer Modelling*, 48(9–10):1628–1645, 2008.
 - [53] IACR Elections. <http://www.iacr.org/elections/> (accessed 7 May 2014), 2013.
 - [54] Markus Jakobsson and Ari Juels. Mix and Match: Secure Function Evaluation via Ciphertexts. In *ASIACRYPT’00: 6th International Conference on the Theory and Application of Cryptology and Information Security*, volume 1976 of *LNCS*, pages 162–177. Springer, 2000.
 - [55] Markus Jakobsson, Ari Juels, and Ronald L. Rivest. Making Mix Nets Robust for Electronic Voting by Randomized Partial Checking. In *11th USENIX Security Symposium*, pages 339–353, 2002.
 - [56] Douglas W. Jones and Barbara Simons. *Broken Ballots: Will Your Vote Count?*, volume 204 of *CSLI Lecture Notes*. Center for the Study of Language and Information, Stanford University, 2012.
 - [57] Ari Juels, Dario Catalano, and Markus Jakobsson. Coercion-Resistant Electronic Elections. Cryptology ePrint Archive, Report 2002/165, 2002.
 - [58] Ari Juels, Dario Catalano, and Markus Jakobsson. Coercion-Resistant Electronic Elections. In *WPES’05: 4th Workshop on Privacy in the Electronic Society*, pages 61–70. ACM Press, 2005. See also <http://www.colombia.rsa.com/rsalabs/staff/bios/ajuels/publications/Coercion/Coercion.pdf> (accessed 7 May 2014).
 - [59] Ari Juels, Dario Catalano, and Markus Jakobsson. Coercion-Resistant Electronic Elections. In David Chaum, Markus Jakobsson, Ronald L. Rivest, and Peter Y. A. Ryan, editors, *Towards Trustworthy Elections: New Directions in Electronic Voting*, volume 6000 of *LNCS*, pages 37–63. Springer, 2010.
 - [60] Aggelos Kiayias. Electronic voting. In *Handbook of Financial Cryptography and Security*, chapter 3. Chapman and Hall/CRC, 2010.
 - [61] Aggelos Kiayias, Thomas Zacharias, and Bingsheng Zhang. End-to-end verifiable elections in the standard model. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology - EUROCRYPT 2015*, volume 9057 of *Lecture Notes in Computer Science*, pages 468–498. Springer Berlin Heidelberg, 2015.
 - [62] Steve Kremer and Mark D. Ryan. Analysis of an Electronic Voting Protocol in the Applied Pi Calculus. In *ESOP’05: 14th European Symposium on Programming*, volume 3444 of *LNCS*, pages 186–200. Springer, 2005.
 - [63] Steve Kremer, Mark D. Ryan, and Ben Smyth. Election verifiability in electronic voting protocols. In *ESORICS’10: 15th European Symposium on Research in Computer Security*, volume 6345 of *LNCS*, pages 389–404. Springer, 2010.
 - [64] Ralf Küsters, Tomasz Truderung, and Andreas Vogt. Accountability: Definition and relationship to verifiability. Cryptology ePrint Archive, Report 2010/236 (version 20150202:163211), 2015.
 - [65] Ralf Küsters. Private email communication, 24th June 2014.
 - [66] Ralf Küsters. Private email communication, October/November 2014.
 - [67] Ralf Küsters and Tomasz Truderung. An Epistemic Approach to Coercion-Resistance for Electronic Voting Protocols. In *S&P’09: 30th IEEE Symposium on Security and Privacy*, pages 251–266. IEEE Computer Society, 2009.
 - [68] Ralf Küsters, Tomasz Truderung, and Andreas Vogt. Accountability: Definition and relationship to verifiability. In *Proceedings of the 17th ACM Conference on Computer and Communications Security*, pages 526–535. ACM, 2010.
 - [69] Ralf Küsters, Tomasz Truderung, and Andreas Vogt. Verifiability, Privacy, and Coercion-Resistance: New Insights from a Case Study. In *S&P’11: 32nd IEEE Symposium on Security and Privacy*, pages 538–553. IEEE Computer Society, 2011. Full version available at <http://infsec.uni-trier.de/publications/paper/KuestersTruderungVogt-TR-2011.pdf>.
 - [70] Ralf Küsters, Tomasz Truderung, and Andreas Vogt. A Game-Based Definition of Coercion-Resistance and its Applications. *Journal of Computer Security*, 20(6):709–764, 2012.
 - [71] Ralf Küsters, Tomasz Truderung, and Andreas Vogt. Clash Attacks on the Verifiability of E-Voting Systems. In *S&P’12: 33rd IEEE Symposium on Security and Privacy*, pages 395–409. IEEE Computer Society, 2012.
 - [72] Tal Moran and Moni Naor. Receipt-Free Universally-Verifiable Voting with Everlasting Privacy. In *CRYPTO’06: 26th International Cryptology Conference*, volume 4117 of *LNCS*, pages 373–392. Springer, 2006.
 - [73] C. Andrew Neff. Practical high certainty intent verification for encrypted votes. Unpublished manuscript, 2004.
 - [74] Helios Princeton Elections. <https://princeton.heliosvoting.org/> (accessed 7 May 2014), 2012.
 - [75] Participants of the Dagstuhl Conference on Frontiers of E-Voting. *Dagstuhl Accord*, 2007. <http://www.dagstuhlaccord.org/> (accessed 7 May 2014).
 - [76] R. A. Peters. A secure bulletin board. Master’s thesis, Technische Universiteit Eindhoven, June 2005.
 - [77] Kazue Sako and Joe Kilian. Secure Voting Using Partially Compatible Homomorphisms. In *CRYPTO’94: 14th International Cryptology Conference*, volume 839 of *LNCS*, pages 411–424. Springer, 1994.
 - [78] Daniel Sandler and Dan S. Wallach. Casting votes in the Auditorium. In *EVT’07: Electronic Voting Technology Workshop*, Berkeley, CA, USA, 2007. USENIX Association.
 - [79] Nicole Schweikardt. Arithmetic, first-order logic, and counting quantifiers. *ACM Trans. Comput. Logic*, 6(3):634–671, July 2005.
 - [80] Ben Smyth. *Formal verification of cryptographic protocols with automated reasoning*. PhD thesis, School of Computer Science, University of Birmingham, 2011.
 - [81] Ben Smyth and David Bernhard. Ballot secrecy and ballot independence coincide. In *ESORICS’13: 18th European Symposium on Research in Computer Security*, volume 8134 of *LNCS*, pages 463–480. Springer, 2013.

- [82] Ben Smyth and David Bernhard. Ballot secrecy with malicious bulletin boards. Technical Report 2014/822, Cryptology ePrint Archive, 2014.
- [83] Ben Smyth, Mark D. Ryan, Steve Kremer, and Mounira Kourjeh. Towards automatic analysis of election verifiability properties. In *ARSPA-WITS'10: Joint Workshop on Automated Reasoning for Security Protocol Analysis and Issues in the Theory of Security*, volume 6186 of *LNCS*, pages 165–182. Springer, 2010.
- [84] *Key issues and conclusions: May 2007 electoral pilot schemes*, May 2007. http://www.electoralcommission.org.uk/_data/assets/electoral_commission_pdf_file/0015/13218/Keyfindingsandrecommendationssummarypaper_27191-20111_E_N_S_W_.pdf (accessed 7 May 2014).
- [85] Filip Zagórski, Richard T. Carback, David Chaum, Jeremy Clark, Aleksander Essex, and Poorvi L. Vora. Remotegrity: Design and use of an end-to-end verifiable remote voting system. In *Applied Cryptography and Network Security*, volume 7954 of *Lecture Notes in Computer Science*, pages 441–457. Springer, 2013.